

ECONOMIA ITALIANA

Fondata da Mario Arcelli

Intelligenza artificiale:
da innovazione tecnologica
a crescita economica
inclusiva

2026/2

 LUISS

CASMEF Centro Arcelli
per gli Studi Monetari e Finanziari

UNIVERSITÀ CATTOLICA del Sacro Cuore
CESPEM

Centro Studi di Politica economica
e monetaria "Mario Arcelli"

Economia Italiana

Fondata da Mario Arcelli

COMITATO SCIENTIFICO

(Editorial board)

CO-EDITORS

FRANCESCO NUCCI - Sapienza, Università di Roma

PAOLA PROFETA - Università Bocconi

PIETRO REICHLIN - Luiss Guido Carli

MEMBRI DEL COMITATO *(Associate Editors)*

CARLOTTA BERTI CERONI

Università di Bologna

MASSIMO BORDIGNON

Università Cattolica del Sacro Cuore

LORENZO CODOGNO

London School of Economics and Political Science

GIUSEPPE DE ARCANGELIS

Sapienza, Università di Roma

STEFANO FANTACONE

Centro Europa Ricerche

EMMA GALLI

Sapienza, Università di Roma

PAOLO GIORDANI

LUISS Guido Carli

GIORGIA GIOVANNETTI

Università di Firenze

ENRICO GIOVANNINI

Università di Roma "Tor Vergata"

ANDREA MONTANINO

Cassa Depositi e Prestiti

SALVATORE NISTICÒ

Sapienza, Università di Roma

ALESSANDRO PANDIMIGLIO

Università degli Studi "Gabriele d'Annunzio" Chieti - Pescara

ALBERTO PETRUCCI

Luiss Guido Carli

GUSTAVO PIGA

Università di Roma "Tor Vergata"

BENIAMINO QUINTIERI

Università di Roma "Tor Vergata"

FABIANO SCHIVARDI

LUISS Guido Carli

MARCO SPALLONE

Università degli Studi "Gabriele d'Annunzio" Chieti - Pescara

FRANCESCO TIMPANO

Università Cattolica del Sacro Cuore

MARIO TIRELLI

Università Roma Tre

GIOVANNA VALLANTI

LUISS Guido Carli

DIRETTORE RESPONSABILE *(Editor in Chief)*: GIOVANNI PARRILLO

ADVISORY BOARD

PRESIDENTE *(President)*

PAOLO GUERRIERI - Sapienza, Università di Roma e PSIA, SciencesPo

CONSIGLIO *(Members)*

FEDERICO ARCELLI, Center for International Governance Innovation

RICCARDO BARBIERI, Tesoro

MARIANO BELLA, Confcommercio

PIERO CIPOLLONE, Banca Centrale Europea

STEFANO COSTA, Istat

GIORGIO DI GIORGIO, Editrice Minerva Bancaria

DARIO FOCARELLI, Ania

RICCARDO GABRIELLI, Deloitte

MAURO MICILLO, Intesa Sanpaolo

STEFANO MICOSI, Assonime

FRANCESCO MINOTTI, "Luigi Luzzatti" s. c. p. a.

ROBERTO MONDUCCI, Osservatorio Imprese Estere

FABIOLA RICCARDINI, Istat

ALBERTO TOSTI, Sara Assicurazioni

MARCO VULPIANI, Deloitte

Economia italiana

Fondata da Mario Arcelli

Intelligenza artificiale:
da innovazione
tecnologica
a crescita economica
inclusiva

numero 2/2026

Pubblicazione quadrimestrale

Roma

ECONOMIA ITALIANA

Rivista quadrimestrale fondata nel 1979 da Mario Arcelli

DIRETTORE RESPONSABILE

Giovanni Parrillo, Editrice Minerva Bancaria

COMITATO DI REDAZIONE

Simona D'Amico (*coordinamento editoriale*)

Francesco Baldi

Nicola Borri

Giuseppina Gianfreda

Rita Mascolo

Guido Traficante

(Pubblicità inferiore al 70%)

Autorizzazione Tribunale di Roma n. 43/1991

ISSN: 0392-775X

Gli articoli firmati o siglati rispecchiano soltanto il pensiero dell'Autore e non impegnano la Direzione della Rivista.

I *saggi* della parte monografica sono a invito o pervengono a seguito di call for papers e sono valutati dall'editor del numero.

I *contributi* vengono valutati anonimamente da due referee individuati dagli editor o dalla direzione e redazione con il supporto dei membri del Comitato Scientifico.

Le *rubriche* sono sottoposte al vaglio della direzione/redazione.

Finito di stampare nel mese di settembre 2026

www.economiaitaliana.org - redazione@economiaitaliana.org

Editrice Minerva Bancaria srl

DIREZIONE E REDAZIONE Largo Luigi Antonelli, 27 – 00145 Roma

AMMINISTRAZIONE

EDITRICE MINERVA BANCARIA S.r.l.

presso Ptsclas, Viale di Villa Massimo, 29

00161 - Roma

amministrazione@editriceminervabancaria.it

Segui Editrice Minerva Bancaria su:



Sommario

Intelligenza artificiale: da innovazione tecnologica a crescita economica inclusiva

EDITORIALE

- 5 Intelligenza artificiale: da innovazione tecnologica a crescita economica inclusiva
A. Bonfiglioli, G. Gancia, P. Profeta

SAGGI

- 13 Intelligenza artificiale e produttività: canali, evidenze empiriche e sfide per l'Italia
F. Calvino, A. Ughi
- 109 Consequences of AI-Driven Chatbots in Economic Domains
S. Galletta
- 153 Workers and Generative AI: Evidence from Italy
M. Anelli, I. Colantone, A. Gallego, P. Stanig
- 227 Exposure to AI, Risk Sharing and Within-Household Inequality
E. Colombo, F. Trentini
- 279 AI and Workplace Welfare Plans
M. Caselli, E. Cristoforetti, A. Fracasso

INTERVENTI

- 323 L'Intelligenza Artificiale nei mercati finanziari italiani:
dall'analisi all'azione
C. Di Noia

RUBRICHE

- 345 Oltre Hormuz. Sicurezza energetica e nuove dipendenze
M. Parco

RECENSIONI

- 357 R. Brunetta (a cura di) *Corpi Intermedi, società civile
ed economia sociale. Per un governo delle transizioni*
M. Rotili

Consequences of AI-Driven Chatbots in Economic Domains

Sergio Galletta *

Abstract

Generative AI has diffused across the economy and public life at exceptional speed. This paper surveys the first wave of empirical evidence on its consequences, focusing primarily on chatbot-style systems. The evidence is clearest for structured workplace tasks, where generative AI tools often raise productivity and reduce performance gaps between novice and expert workers, though aggregate labor-market effects remain limited. In politics and public services, conversational tools lower information costs for voters and citizens but also reduce the cost of persuasion and synthetic content. Findings in education, entrepreneurship, and innovation are more mixed: outcomes depend heavily on how the tools are designed and used, and individual productivity gains can coexist with concerns about collective diversity. Conversational AI also reshapes social behavior, with benefits in brief therapeutic interactions but signs of dependence and weakened accountability under sustained use.

* Sapienza University of Rome, sergio.galletta@uniroma1.it

Sintesi - Conseguenze dei chatbot basati sull'intelligenza artificiale per l'economia

L'intelligenza artificiale generativa si è diffusa nell'economia e nella vita pubblica con una rapidità senza precedenti. Questo articolo passa in rassegna le prime evidenze empiriche sui suoi effetti, concentrandosi in particolare sui chatbot basati su IA generativa. I risultati più consolidati riguardano le attività lavorative strutturate, nelle quali i chatbot tendono ad aumentare la produttività e a ridurre i divari di performance tra lavoratori meno esperti ed esperti, mentre gli effetti aggregati sul mercato del lavoro restano finora limitati. In ambito politico e nei servizi pubblici, questi strumenti riducono i costi di accesso all'informazione per elettori e cittadini, ma abbassano al tempo stesso i costi della produzione di contenuti generati artificialmente. Le evidenze relative a istruzione, imprenditorialità e innovazione sono invece più eterogenee: gli effetti dipendono in larga misura da come i chatbot vengono progettati e utilizzati, e i guadagni di produttività individuale possono accompagnarsi a rischi di omologazione delle idee e dei contenuti. I chatbot incidono infine anche sui comportamenti sociali, mostrando potenziali benefici nelle interazioni terapeutiche di breve periodo, ma anche segnali di dipendenza e di indebolimento del senso di responsabilità individuale quando l'uso si prolunga nel tempo

JEL Classification: O33; J24; J31; D72; D83; H41; L86.

Parole chiave: *Intelligenza artificiale generativa; Chatbot; Agenti di IA; Produttività; Mercati del lavoro; Pubblica amministrazione; Media; Democrazia.*

Keywords: **generative** AI; Chatbots; AI agents; Productivity; Labor markets; Public administration; Media, Democracy.

1. Introduction

In November 2022, OpenAI released ChatGPT. Within two months, it was estimated to have reached 100 million monthly active users, making it the fastest-growing consumer application then recorded. By late 2025, 54.6 percent of U.S. adults aged 18–64 and 32.7 percent of people aged 16–74 in the EU reported using generative AI tools (Bick, Blandin, and Deming, 2025; Eurostat, 2025b); among EU residents aged 16–24, the share was 63.8 percent (Eurostat, 2026). The technology has moved well beyond experimental use. Firms now route customer queries through AI assistants, use AI voice agents for first-round job interviews, and embed generative tools in coding and analytical workflows. Schools deploy AI tutors. Governments use conversational interfaces to help citizens navigate taxes, immigration, and public benefits.¹ The speed of this diffusion has also generated interest among academics, who have contributed a rapidly growing body of rigorous empirical research.

The goal of this paper is to survey that evidence. It covers workplace productivity, hiring, education, public administration, political communication, scientific research, entrepreneurship, and some early findings on labor-market adjustment. The scope is broad because the technology is broad: the same models that can help draft customer replies can also screen job applicants, tutor students, and generate political messages.

The paper begins with evidence on chatbot use in completing specific tasks. The strongest results come from customer support, writing, consulting, and software development, where AI raises productivity and often reduces

¹ For public-sector examples, see the gov.pt assistant and the Canada Revenue Agency's GenAI beta (gov.pt, 2024; Canada Revenue Agency, 2025).

performance gaps within tasks (Brynjolfsson, Li, and Raymond, 2025; Noy and Zhang, 2023; Dell’Acqua et al., 2026a; Cui et al., 2026). Despite this micro-evidence, there is still little sign of large aggregate changes in wages or employment. At this stage, that absence of evidence is probably more informative about slow firm adjustment than about the technology’s ultimate effects. The paper then reviews the growing literature on hiring, education, public administration, and political communication, where the outcomes of interest are different: screening quality, learning, legitimacy, persuasion, and trust. In these domains, the gains in speed and competence coexist with open questions about accountability and the quality of human oversight.

The studies reviewed below rest on a mix of empirical strategies. The strongest causal claims come from experiments in which access to AI is randomized across users, workers, or firms (Noy and Zhang, 2023; Fischer, Rau, and Rilke, 2025; Ash, Galletta, and Opocher, 2025; Ju and Aral, 2025; Brynjolfsson, Li, and Raymond, 2025; Jabarian and Henkel, 2025). A separate body of work compares exposed and unexposed groups before and after the arrival of generative AI, often relying on administrative or near-universe data to track changes in employment, wages, and vacancies (Hui, Reshef, and Zhou, 2024; Johnston and Makridis, 2025; Chen, Srinivasan, and Zakerinia, 2024; Humlum and Vestergaard, 2025; Kauhanen and Rouvinen, 2026; Facius and Iacono, 2026). The paper also draws on descriptive evidence from adoption surveys, platform telemetry, and observational analyses of scientific publications (Bick, Blandin, and Deming, 2026; Chatterji et al., 2025; Anthropic, 2026a; Sarkar, 2026; Hao et al., 2026).

The literature is expanding at an unusual pace, and this survey reflects the state of the evidence as of early 2026. The aim is not to be exhaustive but to identify the main patterns and the questions they raise. A caveat that runs


through what follows is that most of the credible estimates are partial-equilibrium: they measure the effect of giving AI access to a treated worker, student, voter, or firm, holding the rest of the environment fixed.

2. Definitions and diffusion

Generative AI is the broad category used in this paper to describe systems based on large language models (LLMs) that generate text, code, and related outputs in response to natural-language prompts. The technology rests on the transformer architecture introduced by Vaswani et al. (2017), scaled up through massive pretraining on text data and refined through reinforcement learning from human feedback. For the purposes of this survey, however, the most important distinction is not between model families but between modes of use.

The first is the chatbot: a conversational system in which a user asks a question or gives an instruction, the model produces a response, and the human decides what to do next. ChatGPT, Claude, Gemini, and similar systems are most often used in this way. Most of the empirical evidence reviewed in this paper concerns this chatbot mode. The second is the AI agent. Unlike a chatbot, an agent can break a task into steps, use external tools, and act across multiple stages under relatively limited human supervision (OECD, 2026). That distinction matters because the economic questions are not identical. The chatbot literature mostly studies assistance within tasks; agentic systems raise additional issues about reliability, monitoring, and control. Since rigorous research on agents remains limited, this paper treats them mainly as a frontier for future research rather than as the core object of the survey.

Figure 1 Public attention to AI-related terms, 2020–2026



These tools quickly entered public attention and everyday use. For descriptive purposes, Figure 1 illustrates both the demand and supply sides of public attention. Panel (a) tracks search interest: “ChatGPT” spiked immediately on release in November 2022 and has remained high, while “LLM” and “AI chatbot” rose more gradually through 2023 and 2024 as the technology entered workplaces and public services. “Agentic AI” was virtually absent until

mid-2024 and has climbed steeply since, reflecting the broadening of the technology's scope from conversation to autonomous action. Panel (b) tracks media coverage using the GDELT Project's monitoring of English-language online news.² The patterns are broadly similar but not identical: news coverage of "ChatGPT" has been more persistent than search interest, and the newer terms appear in news slightly before they enter mainstream queries.

The evidence on attention is consistent with the scale of diffusion reported in surveys. Eurostat reports that 32.7 percent of people aged 16–74 in the EU used generative AI tools in 2025, rising to 63.8 percent among those aged 16–24 (Eurostat, 2025b, 2026). In the United States, the Real-Time Population Survey places overall adoption among adults aged 18–64 at 54.6 percent by August 2025, while 37.4 percent of employed respondents reported using generative AI for work (Bick, Blandin, and Deming, 2025). Gallup's quarterly tracking of U.S. employees suggests that the intensive margin is rising even as the extensive margin flattened in late 2025: 26 percent of workers reported using AI at least a few times per week, up from 23 percent a quarter earlier, while 12 percent used it daily. Use is concentrated in leadership roles: frequent use reached 44 percent among leaders and 30 percent among managers, compared with 23 percent among individual contributors. Technology, finance, and higher education led in adoption, while retail and manufacturing lagged behind (Gallup, 2026). Although these surveys are not directly comparable, organization-level adoption appears slower than individual use: 20 percent of EU firms with at least ten employees reported using AI technologies in 2025, and ECB survey evidence suggests that employee use often runs ahead of dedicated firm-level AI investment (Eurostat, 2025a; European Central Bank,

2 GDELT (the Global Database of Events, Language, and Tone) is an open database that monitors broadcast, print, and web news from around the world.

2026). Similarly, platform telemetry shows substantial use of the technology. OpenAI reports more than one million business customers, and Anthropic's Economic Index documents how usage mixes differ across Claude.ai, enterprise, and API contexts (OpenAI, 2025; Anthropic, 2026a). Governments are also adopting front-end conversational tools, largely for navigation and service rather than unconstrained adjudication. Examples include gov.pt, gov.gr, GOV.UK Chat, the OECD's mAigov case, and the Canada Revenue Agency's beta chatbot (gov.pt, 2024; gov.gr, 2026; GOV.UK, 2025; OECD, 2025; Canada Revenue Agency, 2025).

What do users themselves want from the technology? Anthropic's large-scale qualitative study of 80,508 Claude users across 159 countries, conducted through AI-led open-ended interviews, offers one of the broadest demand-side pictures to date (Anthropic, 2026b). The most common aspirations include professional excellence, personal transformation, and life management, while the most common concerns are unreliability, jobs and the economy, and loss of human autonomy and agency. An important finding is that hopes and concerns often coexist within the same individuals rather than sorting users into clearly pro-AI and anti-AI camps. The study also reports more positive sentiment toward AI in lower- and middle-income countries, where respondents more often emphasize opportunity, entrepreneurship, and access to education, while respondents in wealthier regions more often frame AI as a tool for managing complexity. The study is descriptive and limited to users of a single platform, but its scale makes it a useful complement to standard representative survey data.³ A complementary perspective comes from

3 Complementary evidence based on direct data on ChatGPT conversations shows that most consumer use is non-work-related rather than tied to paid work (Chatterji et al., 2025). In particular, they find that non-work-related messages account for more than 70 percent of usage, and that "Practical Guidance", "Seeking Information", and "Writing" together make up nearly 80 percent of all conversations. This suggests that a substantial share of the gains from these tools may take the form of home production and reduced decision costs in every-

more structured experimental evidence on stated demand. In a preregistered conjoint experiment with a representative U.S. sample of nearly 2,000 respondents, Ash, Capozza, and Galletta (2026) estimate willingness to pay for five categories of AI conversational apps: administrative assistance, study and tutoring, wellness coaching, friendship, and romantic companionship. Willingness to pay is highest for administrative assistance and study/tutoring, and substantially lower for relational uses such as friendship and romantic companionship.

3. Productivity, work, and entrepreneurship

When a new technology enters the economy, the first question economists typically ask is whether it raises productivity. Generative AI has been no exception: much of the early empirical literature studies settings in which tasks are well defined, performance is observable, and access to the tool can be randomized or quasi-experimentally assigned. For example, Brynjolfsson, Li, and Raymond (2025) study the staggered rollout of a generative-AI assistant to customer support workers at a large firm. They find that access to the system raised issues resolved per hour by about 15 percent on average, with especially large gains for less experienced and lower-skill workers. The authors argue that the tool helped disseminate the best practices of stronger workers and allowed newer workers to move more quickly down the experience curve. The result was both a level effect on output and a compression of performance gaps within the firm.

day life, which are difficult to measure in standard labor-market outcomes.

Similar results appear in other studies. In a preregistered online experiment, Noy and Zhang (2023) find that access to ChatGPT reduced completion time for professional writing tasks by about 40 percent while improving quality by about 18 percent. In a field experiment with management consultants at BCG, Dell’Acqua et al. (2026a) show that GPT-4 access improved speed and output quality for tasks well inside the model’s capability frontier,⁴ but reduced accuracy when workers relied on it for a task outside that frontier. The study highlights both sides of the technology: large returns within the frontier and the risk of degraded performance outside it. In separate follow-up work, Dell’Acqua et al. (2026b) examine human-AI teams in a collaborative production setting, finding that individuals working with AI performed about as well as two-person teams without it, and that AI reduced functional silos across professional groups by helping workers produce more balanced solutions regardless of background.

Additional credible field results come from software development. Cui et al. (2026) report three randomized controlled trials at Microsoft, Accenture, and an anonymous Fortune 100 firm, covering 4,867 developers, with a random subset assigned access to GitHub Copilot, a widely used coding assistant. Because compliance was imperfect, the authors use treatment assignment as an instrument for actual usage. Their preferred instrumental-variable estimate implies a 26 percent increase in weekly completed tasks, with less experienced developers showing both higher adoption rates and larger productivity gains. Further experimental results point in the same direction. Cruces et al. (2026) conduct a randomized experiment with 1,174 adults completing an incentivized workplace-style business problem-solving task and show that AI

4 In this context, the model’s “capability frontier” means the range of tasks for which the system can produce reliable, decision-useful output with limited correction. Inside that range, the model helps with execution; outside it, the output remains fluent but becomes harder to trust.

access closes roughly three-quarters of the education-based productivity gap, reducing it from 0.548 to 0.139 standard deviations.

The studies discussed so far measure productivity on well-defined tasks. A second strand of work moves beyond the single-task frame and looks at how generative AI reshapes broader organizational margins: how workers allocate their time, how firms hire, and how decision-making unfolds in open-ended environments.

First, generative AI appears to change how knowledge workers organize their day. In a six-month randomized field experiment across 66 firms with 7,137 knowledge workers, Dillon et al. (2025) study the effects of access to a generative AI tool integrated into email, document, and meeting applications. Active users spent less time on email — about two fewer hours per week in the second half of the experiment — and completed documents moderately faster, but the authors do not detect large shifts in the quantity or composition of workers' tasks. The result captures a margin that task-level experiments often miss: not whether a single memo improves, but whether the overall structure of the workday shifts when AI is available for months rather than minutes.

Second, hiring is another organizational margin that can affect firm productivity. Jabarian and Henkel (2025) conduct a field experiment covering roughly 70,000 applicants to compare human recruiters with AI voice agents conducting automated job interviews. Applicants interviewed by AI were more likely to receive offers, more likely to start jobs, and more likely to remain after 30 days, with no decline in post-hire productivity. When allowed to choose, a large majority of applicants selected the AI option. Transcript analysis suggests that the gain comes from structure: AI interviews follow guidelines more consistently, cover more hiring-relevant topics, and produce a more standardized information set for downstream evaluation.

Finally, decision-making in entrepreneurship shows a different pattern. In a five-month randomized field experiment with 640 Kenyan small business entrepreneurs, half of whom were assigned access to a GPT-4-powered AI business assistant over WhatsApp, Otis et al. (2026) find no average treatment effect on business performance. But that average masks heterogeneity: lower-performing entrepreneurs at baseline did about 8 percent worse, while higher-performing entrepreneurs may have benefited by just over 15 percent. The paper finds no evidence that this difference is driven by the quantity or type of questions asked, or by systematic differences in the advice generated. Instead, the authors' preferred interpretation is that heterogeneity arises from which pieces of AI advice entrepreneurs selected and implemented: higher performers appear more likely to adopt tailored, context-specific suggestions, whereas lower performers more often adopt generic recommendations such as discounting and advertising. This suggests that when AI is used in open-ended settings, outcomes may depend less on the average quality of advice than on users' judgment in selecting and acting on it.

Whether such task-level gains translate into aggregate outcomes is a harder question, and the existing micro studies are not well placed to answer it. They identify the effect of generative AI on particular tasks, workers, or firms, while the aggregate object depends on reallocation across tasks, the creation of new ones, and changes in labor supply. The task-based framework of Acemoglu and Restrepo (2018, 2019) provides a useful discipline: a technology can raise productivity within a task without implying mechanical changes in wages, employment, labor shares, or welfare. Acemoglu (2025) applies this logic to AI and argues that, even under generous assumptions about exposure and within-task gains, the implied effect on total factor productivity over a ten-year horizon is modest. A 15 percent improvement on a subset of tasks

therefore does not mechanically imply a comparable reduction in labor demand: firms may expand output, reallocate workers toward complementary activities, or devote more time to monitoring and verification. The relevant question is whether those gains are large and broad enough to affect earnings, hours, and employment in the aggregate, and on that margin the evidence remains early and mixed.

Early null results come from Denmark. Humlum and Vestergaard (2025) link large-scale adoption surveys to administrative labor records covering earnings, hours, and job histories. Using a difference-in-differences design, they estimate near-zero effects on earnings and recorded hours two years after ChatGPT's introduction, with confidence intervals that rule out effects larger than 2 percent. The null holds for intensive users, early adopters, workplaces with substantial AI investments, and occupations with flexible pay. This is informative because Denmark's labor market, with low firing costs, decentralized wage bargaining, and annual pay renegotiation, is one where adjustment to a productivity shock should appear relatively quickly if it is going to appear at all. A Finnish study using similar administrative data finds no comparable displacement either (Kauhanen and Rouvinen, 2026), and Facius and Iacono (2026) reach the same conclusion in Norway using population-wide registers through early 2025 and three complementary designs; a backdating exercise shows that the small negative point estimates among young workers in AI-exposed occupations are absorbed by pre-existing trends.

Evidence from other settings, however, points to heterogeneous adjustment rather than a uniform null. In highly exposed segments of the online labor market, freelance demand falls noticeably (Hui, Reshef, and Zhou, 2024). U.S. sectoral findings are mixed: Johnston and Makridis (2025) find output and employment gains in some AI-exposed sectors alongside a lower labor

share, while payroll data show pressure on early-career workers in exposed occupations (Brynjolfsson, Chandar, and Chen, 2025). Cross-country EU studies are also heterogeneous: employment gains are stronger in AI-exposed occupations, but appear to depend on complementary innovation capacity (Guarascio and Reljic, 2025). Across these studies, adjustment is concentrated in particular sub-markets and worker groups rather than spread evenly across the economy.

Firm-side surveys tell a complementary story. Baslandze et al. (2026) survey nearly 750 corporate executives through the Atlanta Fed–Richmond Fed–Duke CFO Survey and find that more than half of firms have already invested in AI, with productivity gains concentrated in high-skill services and finance. Yet half of the responding firms explicitly report that AI will not replace any roles, and the main motivation for investment is improving efficiency rather than cutting labor costs. Executives also perceive larger productivity gains than are implied by observed changes in revenue and employment, a wedge that likely reflects delayed revenue realization and quality improvements not yet captured in measured output.

Two recent papers help interpret why the labor-market effects are so uneven. Using a near-universe panel of U.S. job postings, Chen, Srinivasan, and Zakerinia (2024) distinguish occupations in which generative AI is more likely to automate structured cognitive tasks from those in which it is more likely to augment human work: vacancy demand falls in the first group and rises in the second. Althoff and Reichardt (2026) push the argument further in a dynamic task model whose distinctive margin is simplification — AI reduces the skill threshold required to perform some tasks well. That idea reconciles two otherwise awkward facts: within a task, AI compresses performance differences by making good practice easier to access; across jobs, the same sim-

plification puts pressure on entry margins and changes the value of expertise.

The task-level estimates are real, but their translation into aggregate labor-market outcomes depends on firm reorganization, workflow redesign, and the pace at which complementary capital is built, most of which current micro studies cannot observe directly.

4. Learning and human capital

In most workplace settings, the relevant measure of success is current output: how many tickets were resolved, how fast the memo was drafted. In education, current output matters too, as students need to pass exams and complete assignments, but it is only valuable insofar as it reflects durable knowledge. A high grade produced with AI assistance has limited human-capital value if the student cannot reproduce the reasoning without help. That gap between assisted performance and retained learning makes education one of the most revealing settings in which to study generative AI.

The central question is therefore not whether AI raises performance during practice — by now there is broad support that it does — but whether students retain that performance once the tool is taken away. The studies reviewed below address both sides of that question, and the answer depends heavily on how the AI is used.

Some of the findings are encouraging. Tutoring studies point to real gains when AI is embedded in a learning design rather than treated as a pure substitute for student effort. Fischer, Rau, and Rilke (2025) conduct a randomized experiment with 334 university students preparing for an incentivized exam. Students were assigned to a control group, a restricted-access AI tutor, or un-

restricted AI tutor access during a 25-minute learning period, and were then tested without the tool. Overall AI tutor access raised unaided test performance by 0.23 standard deviations relative to control. The result was driven by unrestricted access, which outperformed the control group by 0.34 standard deviations and the restricted-access group by 0.21 standard deviations. Importantly, the behavioral evidence suggests that students still engaged with the underlying material rather than using the tutor to avoid reading effort. In a complementary classroom study, Vanzo, Chowdhury, and Sachan (2025) replace traditional homework with GPT-4-based interactive sessions in four Italian high-school classes learning English as a second language. Using a prompting strategy that requires minimal content preparation, they report greater gains in grammar, higher student engagement, and high student satisfaction. These studies suggest that conversational AI can improve learning when it is used to provide practice, explanation, and feedback rather than simply to generate answers.

But other studies find the opposite: AI use can erode learning rather than support it. A student who finishes today's assignment with AI may or may not know more tomorrow without it. Bastani et al. (2025) provide direct evidence on this distinction in a field experiment with nearly 1,000 high-school math students in Turkey. They compare a standard GPT-4-style interface, a tutor-style interface with guardrails that guide problem solving rather than provide direct answers, and a control group using only traditional materials. The unrestricted AI group performed well during AI-assisted practice but did worse on subsequent unaided exams than the control group, whereas the guardrailed tutor group did not suffer the same penalty. Strömberg, Lei, and Wu (2026) document the same homework-versus-exam wedge at much larger scale in China: in 30-month panel data on 26,811 secondary-school students,

AI adoption raises homework scores by about 18 percent but lowers monthly exam scores by 20 percent within six months, and high-stakes entrance-exam scores fall by 18 to 24 percent over about two years. The losses are concentrated among roughly 80 percent of users whose pattern of short completion times and high homework scores is consistent with homework outsourcing.

Evidence from higher education is more mixed. Hausman, Rigbi, and Weisburd (2025) use student-course data covering roughly 36,000 students in 6,000 courses at a leading Israeli university, comparing AI-compatible courses with AI-incompatible ones before and after ChatGPT's release. AI availability raised grades in AI-compatible courses by about 1 point at first, rising to 1.5 points in the following year, with larger gains for students at the bottom of the distribution and a one-third reduction in failure rates. But students first exposed to AI in introductory courses did no better, and sometimes worse, in later advanced exam-based courses where AI could not be used; at the same time, they performed better in later AI-compatible courses. The pattern suggests accumulation of AI-specific skills alongside weaker accumulation of traditional human capital, a distinct margin from the substitution-for-effort harms above.

The pattern across these studies is that access alone is not enough; what matters is whether the tool is built to support learning rather than replace effort. If the goal is durable learning, schools and policymakers need purpose-built tools rather than raw chatbots. Chung et al. (2026) illustrate what that can look like. In collaboration with the Taipei City Government, they integrate a reinforcement learning algorithm into a GenAI tutoring platform and deploy it across ten high schools in a five-month Python programming course. The algorithm uses LLM-extracted signals from student-chatbot interactions to estimate each student's knowledge state and adjust problem difficulty in real

time. In a randomized trial with 770 students, adaptive sequencing raised unassisted final exam performance by 0.15 standard deviations — equivalent to roughly six to nine months of additional schooling — without increasing instruction time or teacher workload.

A different line of evidence places the LLM in the hands of a human tutor rather than the student. Wang et al. (2024) study Tutor CoPilot, a system that provides real-time, expert-like suggestions to tutors during one-on-one on-line sessions, in a randomized trial with 900 tutors and roughly 1,800 K–12 students in underserved U.S. communities. Students whose tutors had access to Tutor CoPilot were 4 percentage points more likely to master a topic, a 9 percent gain, and the effect was largest for students assigned to lower-rated tutors, where mastery rose by 9 percentage points. Transcript analysis indicates that the system pushed tutors toward higher-quality pedagogical strategies, such as guiding students to think through problems rather than supplying answers directly. The largest gains in this domain may therefore come not from giving the model directly to the learner, but from using it to upgrade the practice of a relatively inexperienced human tutor. Kestin et al. (2025) provide a controlled comparison in higher education. In a preregistered within-subjects crossover study with fewer than 200 Harvard undergraduates in introductory physics, students learned more than twice as much in less time when using a purpose-built AI tutor at home than when learning the same material in class under best-practice active-learning instruction. The contrast is striking because the comparison condition already represents a refined version of the pedagogy that the AI is being measured against.

Taken together, these results suggest a consistent distinction: when AI substitutes for student effort, immediate performance may improve while learning can suffer; when it is used to guide practice, sequence tasks, and provide

feedback without replacing problem-solving, the implications for human capital are more favorable.

These estimates are short-run and partial-equilibrium: they capture the effect of giving students access to AI under current curricula, exam formats, and grading rules. Universities are already redesigning assessments and honor codes in response, and those adjustments will change what such estimates mean in the medium run.⁵

Finally, the pedagogical question has a strategic dimension. As schools and education authorities increasingly make AI tools available or even encourage their use, parents become key gatekeepers in deciding how much and how freely their children engage with these technologies. Yet even parents who worry about AI's effect on learning may feel they cannot restrict its use if their child's peers are already using it. Bursztyn et al. (2026) study this dynamic using incentivized experiments on parental demand for AI-based learning tools. Willingness to pay for an AI subscription rises by more than 60 percent when parents learn that other teenagers are using the technology. Informing parents about the potential harms of unstructured AI use makes them more pessimistic about AI's educational value, but does not reduce individual demand. Instead, it increases support for banning AI in schools. This creates a coordination problem: parents adopt AI for their children while simultaneously believing that everyone would be better off with tighter restrictions in educational settings. The pattern resembles familiar competitive dynamics in education, such as tutoring arms races and credential inflation, but it is more

5 Princeton's faculty voted in May 2026 to proctor all in-person exams, ending a 133-year-old unproctored honor-code tradition, with AI-assisted cheating cited as a primary motivation. University College London's law faculty now requires every module to include secure assessment formats designed to safeguard against substituting generative AI for the skills being evaluated (Veale et al., 2025). Oral examinations have returned to undergraduate courses at institutions such as Cornell, where instructors describe oral defenses as a way to test whether students understand work that may have been completed with AI assistance.

pronounced here because the technology may directly undermine the learning it is supposed to support.

5. Politics, media, and the public sector

The previous sections suggest that generative AI can raise productivity by making relevant information easier to access, but also that it can erode learning and perform poorly when users lack judgment or oversight, or when the tools are poorly designed. Both sides are relevant in politics and media. On the one hand, political information is complex and costly to absorb, while voters are busy; a well-designed chatbot can lower the cost of asking a specific question and receiving a tailored explanation grounded in relevant sources. On the other hand, the same technology that makes information cheaper also makes persuasion, manipulation, and synthetic content cheaper.

Ash, Galletta, and Opocher (2025) provide experimental evidence on the first margin. In a preregistered three-wave randomized experiment with 1,463 California voters around the November 2024 California ballot propositions, they compare access to BallotBot, a custom chatbot grounded in the official voter guide via retrieval-augmented generation, with access to the voter guide itself. Much like the studies reviewed above, the effects are strongest for those with the least prior knowledge: BallotBot improved accuracy on in-depth questions about ballot propositions by about 18 percent and cut response time by about 10 percent, with the largest gains among initially less-informed voters; on simpler factual questions, where most voters were already accurate, the chatbot had no detectable effect. Moreover, using an incentive-compatible price-elicitation mechanism, the authors show that BallotBot lowered the

perceived cost of acquiring additional political information, again particularly for low-knowledge respondents. Participants assigned to the chatbot were also more engaged: they spent more time gathering information and were more likely to search for information on other propositions. In other words, the tool did not merely help voters find answers to specific questions; it made them more willing to seek out information in the first place. At the same time, the paper shows where the gains stop: in a second wave conducted without access to the tool, the initial assignment to BallotBot had no effect on accuracy or reasoning quality, suggesting that the advantage reflects on-demand competence rather than durable political learning. Finally, BallotBot did not produce detectable effects on self-reported turnout or the direction of voting, although the null effects may partly reflect limited statistical power. Consistent with this, Velez, Green, and Sevi (2025) show in a randomized experiment that a chatbot voting-advice application can inform young politically unaffiliated adults but has only weak downstream effects on party evaluations and vote preferences. These findings suggest that conversational AI may be most effective at reducing search costs and improving comprehension, while having more limited effects on final political decisions.

The other side of the coin is that the same technologies seem to reduce the cost of persuasion. Bai et al. (2025) conduct three preregistered survey experiments with 4,829 U.S. participants, randomly assigning them to read LLM-generated messages, human-written messages, or neutral control texts on polarized policy issues including a smoking ban, an assault weapons ban, and a carbon tax. LLM-generated messages shifted policy attitudes by 2 to 4 points on a 101-point scale — small but consistent effects, comparable in magnitude to those of messages written by lay humans. Interestingly, the persuasion operates through different channels: LLM messages were perceived as more factual,

logical, and dispassionate, while human messages persuaded through perceived originality and authenticity. Salvi et al. (2025) move from static messages to real-time interaction. In a preregistered experiment with 900 participants, they set up short multi-round debates on U.S. political and social issues using a $2 \times 2 \times 3$ factorial design that varies the opponent (human or GPT-4), whether the opponent has access to the participant's sociodemographic profile, and the strength of prior opinion on the topic. In debate pairs where the two sides were not equally persuasive, GPT-4 with personalization won 64.4 percent of the time — an 81 percent relative increase in the odds of shifting the opponent's position. The result shows that conversational AI can use personal information to tailor arguments in ways that can outperform human debaters.

There is also work from real electoral contexts. Lin et al. (2025) conduct preregistered experiments during the 2024 U.S. presidential election, the 2025 Canadian federal election, and the 2025 Polish presidential election, randomizing participants to have conversations with AI models advocating for specific candidates. Across all three countries, they find significant effects on candidate preference, larger than those typically observed for traditional video advertisements. The AI persuades primarily through relevant facts and evidence rather than through sophisticated psychological techniques, but not all of the facts presented were accurate: models advocating for right-leaning candidates consistently produced more inaccurate claims. The paper also reports persistence at follow-up. In parallel, Hackenburg et al. (2025) provide a systematic assessment of what drives AI persuasiveness. Across three experiments with 76,977 U.K. participants, they deploy 19 LLMs on 707 political issues and fact-check 466,769 resulting claims using professional human fact-checkers. They find that model scale delivers only modest returns to persuasiveness, whereas post-training and prompting methods boost it by

as much as 51 and 27 percent, respectively. Critically, where these methods increased persuasiveness, they also systematically decreased factual accuracy. Taken together, these findings suggest that the barrier to scalable political influence may be lower than a focus on frontier-model scale alone would imply.

This body of work shows that both one-shot and interactive political persuasion are becoming cheaper, more personalized, and more scalable. The policy concern is that once persuasion can be generated cheaply, adapted to individual profiles, and deployed at scale, the traditional economics of campaign communication may change substantially. The technology that helps a voter understand a complex ballot proposition also allows campaigns and platforms to target attention and persuasion at very low marginal cost. Which effect dominates depends on the institutional environment.⁶

At the same time, the same technology can also be used in the opposite direction, to correct misinformation, weaken conspiracy beliefs, or reduce inter-group hostility. Costello, Pennycook, and Rand (2024) run two preregistered experiments with more than 2,000 self-identified conspiracy believers in which participants engage in personalized, evidence-based dialogues with GPT-4. A single roughly eight-minute conversation reduces belief in the discussed conspiracy by about 20 percent on average, and the effect persists at two-month follow-up across both political and non-political conspiracies, including conspiracies the participant had not originally discussed. Interactive AI can therefore durably move beliefs that have proven resistant to standard interventions by engaging with the specific reasoning of the believer rather than delivering one-size-fits-all corrections. Piccardi et al. (2025) study a different lever, the algorithmic ranking of political content. In a ten-day

6 For example, Acemoglu, Ozdaglar, and Siderius (2026) formalize how two platform mechanisms, AI-driven recommendations that maximize engagement and targeted digital ads, can create echo chambers, polarize voters, and in turn increase party polarization.

preregistered field experiment with 1,256 X (Twitter) users during the 2024 U.S. presidential campaign, they use an LLM to downrank or uprank content expressing antidemocratic attitudes and partisan animosity. Reducing exposure lowers out-party affective polarization by about two points on a 100-point feeling thermometer, while upranking the same content shifts it in the opposite direction. Platform-level uses of LLMs as moderation infrastructure can therefore have measurable effects on political attitudes that operate through changes in exposure rather than through persuasion. The picture is not uniformly favorable, however. DeVerna et al. (2024) show in a preregistered randomized experiment that exposing participants to fact-checks generated by a large language model does not improve their ability to discern accurate from inaccurate political news, even though the model labels headlines correctly about 90 percent of the time. Two channels drive the harm. When the model expresses uncertainty or mislabels a true headline as false, exposure reduces belief in accurate news. When participants voluntarily access LLM fact-checks, they become significantly more likely to share both true and false news. What these papers suggest is that LLMs in political information environments can correct or amplify distortions depending on how they are designed and deployed, and that institutional choices about platform rules, disclosure, and who controls the prompt may matter at least as much as the model itself.

The information ecosystem surrounding these risks is itself being reshaped. In a field experiment with a major German newspaper, Campante et al. (2025) show that making the threat of AI-generated misinformation salient raises concern about misinformation, reduces trust in news overall, including in the outlet itself, and yet increases visits to the trusted outlet and raises subscriber retention. Cheap synthetic content can erode baseline trust while simultane-

ously increasing the value of credible intermediaries. The implication is that generative AI may not simply weaken the role of trusted media; it may raise the premium on provenance, editorial reputation, and verification.

Finally, the same tension between access and accountability extends to government services. Public administration is in many ways a natural home for chatbots: governments spend enormous resources on activities that are fundamentally conversational, such as explaining rules, routing requests, and helping citizens navigate complex programs. Deployments are already widespread, including Portugal's gov.pt assistant, Greece's mAigov digital assistant, GOV.UK Chat in private beta, and the Canada Revenue Agency's GenAI chatbot beta. In this domain, however, rigorous research on citizen outcomes remains scarce. Existing studies document that government chatbots can improve perceived access and efficiency (Larsen and Folstad, 2024), that citizen adoption depends on perceived usefulness and trust (Alishani and Homburg, 2025), and that public trust in AI-assisted services is sensitive to whether humans are described as remaining in the decision loop (Aoki, 2020, 2021). These findings are drawn from survey experiments and qualitative work rather than from field experiments measuring take-up, resolution times, or administrative burden — a gap that stands in contrast to the richer body of research available in the other domains reviewed in this paper.

The equilibrium consequences in this domain will again depend on how platforms, regulators, and political actors adjust once these tools are deployed at scale. Persuasion estimates are measured holding moderation rules, disclosure requirements, and counter-messaging fixed, and the corrective effects of Costello, Pennycook, and Rand (2024) and Piccardi et al. (2025) are estimated under researcher control rather than in engagement-maximizing pipelines. The gap is widest in the public sector, where most existing evidence concerns

stated trust and perceived usefulness rather than take-up, resolution times, or error rates by group.

6. Innovation, science, and the production of ideas

Another important domain in which to evaluate generative AI is the production of new ideas. From a growth-accounting perspective, what matters is not only whether AI raises the productivity of routine cognitive work, but whether it changes the rate and direction of innovation, research, and creative output. Early findings here point to a sharper version of a pattern already visible in the previous sections: substantial gains at the individual level coexist with concerns about diversity and exploration at the collective level.

Doshi and Hauser (2024) run an online experiment with 293 writers and 600 independent evaluators in which participants are asked to write short stories with or without access to GPT-4 plot ideas. Access to the model raises measured creativity at the individual story level: AI-assisted stories are judged more creative, better written, and more enjoyable on average, and the largest gains accrue to less creative writers, who close part of the gap with their more creative peers. At the same time, the semantic distance between stories within the AI-assisted condition falls, with cosine similarity rising by roughly 9 to 11 percent. The population of stories produced with AI is more homogeneous than the population produced without it. The same model that raises the floor on individual creativity also compresses variation across writers, providing a clean micro-foundation for a concern that recurs throughout this literature: AI can raise individual productivity while reducing collective diversity.

A similar pattern shows up in science. Hao et al. (2026) combine adoption

of AI tools with large-scale bibliometric outcomes. Scientists who adopt AI tools see substantial increases in their personal impact: more papers, more citations, and earlier transitions to project leadership. But the topical breadth of the literature contracts. AI adoption is associated with greater concentration of work on a narrower set of topics, fewer new research lines, and reduced exploration of less popular areas. The aggregate picture is one in which AI accelerates the production of papers within already well-trafficked questions while reducing the discovery rate of new ones. Agrawal, McHale, and Oettl (2026) provide a conceptual framework that fits these patterns: AI in science is best understood as a change in the cost structure of research production rather than as a substitute for scientists, with potentially large effects on which projects get done and which are abandoned.

Two implications follow. The standard productivity narrative extends to the production of ideas, but more strongly for impact than for variety. Individual-level evaluations of AI's effect on creativity or research output are likely to overstate aggregate welfare effects if they ignore the consequences for the diversity of ideas and the long-run rate of discovery. The equilibrium consequences here will also depend on how journals, funders, and academic labor markets reward AI-assisted output, and on whether the observed contraction in topical breadth turns out to be a transient feature of the diffusion process or a structural feature of the technology.

7. Chatbots and social behavior

The evidence reviewed so far treats chatbots primarily as tools for completing tasks, acquiring information, or accessing services. Conversational AI

also changes how people behave and relate to others in ways that go beyond the immediate task. This section reviews experimental evidence on three margins: chatbots as behavioral interventions, chatbots and social connection, and chatbots and moral behavior.

A first set of results shows that chatbot conversations can change behavior at scale. Chopra et al. (2026) conduct a randomized experiment with 2,719 social media users, assigning participants to a control conversation or to one of three AI-delivered intervention protocols: two based on motivational interviewing, a therapeutic technique in which the interviewer guides the participant toward self-persuasion, and one based on direct persuasion. All conversations were conducted by an AI interviewer, enabling standardized delivery at scale. All three interventions increased motivation to change and the perceived costs of social media use, with change-focused self-persuasion yielding the largest effects. The effects persisted and translated into self-reported reductions in social media use more than two weeks later. The result matters because it shows that conversational AI can serve as a delivery mechanism for behavioral interventions that would otherwise require trained human counselors.

A second margin concerns the effect of chatbot use on social connection. The results here are mixed. In a series of studies, De Freitas et al. (2026) show that brief interactions with an AI companion can reduce loneliness. In one experiment, a short conversation with a caring AI companion reduced loneliness about as much as a conversation with another human being. The mechanism is what the authors call “feeling heard” — the perception that another entity listened with attention and empathy. Participants who knew they were talking to a bot benefited much like those who did not. A longer-horizon study, however, points to a more qualified picture. Fang et al. (2025)

randomize 981 participants into a four-week chatbot interaction study with more than 300,000 messages. Higher daily usage was associated with greater loneliness, stronger emotional dependence on the AI, and less socialization with real people. Voice-based interactions appeared somewhat more beneficial at first, but these advantages weakened at higher levels of use. Taken together, the studies suggest a dose-response pattern: brief chatbot interactions can provide genuine support, but sustained heavy use may begin to substitute for rather than complement human relationships.

A third finding concerns the moral consequences of delegating decisions to AI. Köbis et al. (2025), in 13 experiments across four main studies (those involving human participants were preregistered), show that delegation to AI can sharply increase dishonest behavior. In die-roll tasks, honesty was high when participants reported outcomes directly, but requests for cheating increased when participants could induce a machine to behave dishonestly through supervised learning or high-level goal setting. In natural-language delegation tasks, machine agents were also far more likely than human agents to comply with fully unethical instructions. The mechanism is moral distancing: the presence of an automated intermediary weakens the psychological cost of dishonesty.⁷ The result is relevant for organizational settings where AI agents act on behalf of employees or managers, and it connects to the public-administration evidence discussed earlier, where citizens blame AI systems less than government agencies for service failures (Liang, Li, and Chen, 2025) — in both cases, responsibility is diffused through the technology.

These results suggest that the welfare effects of chatbots go beyond the task-level productivity gains that dominate most of the current literature.

7 Complementary evidence from a simpler reporting experiment points in a similar direction. Biener and Waeber (2024) show that people are more honest with human counterparts than with chatbots when agency cues are present.

Chatbots can deliver behavioral and therapeutic conversations at scale and provide emotional support in brief interactions, but they can also foster dependence, reduce real-world socialization, and weaken moral accountability when used as intermediaries. These are margins that standard analyses of AI adoption are unlikely to capture, and in the absence of longitudinal evidence at the population level the findings reviewed here are best read as evidence on mechanisms rather than as estimates of long-run social consequences.

8. Discussion and conclusion

This paper has reviewed the first wave of empirical evidence on generative AI across several economic domains. Three broad patterns emerge. First, results are strongest in structured tasks (e.g., customer support, writing, coding, consulting, and some forms of tutoring) where chatbots often raise performance and compress gaps between lower- and higher-performing users. Second, these gains are highly sensitive to context: outcomes depend on whether the task is easy to verify, whether users can exercise judgment, and whether the tool is designed to support rather than replace human effort. Third, the same technology that lowers information and service costs can also create new risks (e.g., in learning, persuasion, accountability, and social behavior) that do not appear in narrow productivity measures.

A recurring result across domains is that average effects often conceal substantial heterogeneity. Less experienced workers, less-informed voters, and students who receive structured support often benefit the most. But in more open-ended environments, the same technology can amplify differences in judgment, preparation, and institutional setting. In entrepreneurship, for ex-

ample, AI appears to help users who can formulate actionable problems and screen advice well, while harming those who cannot. In education, AI can improve performance when it structures practice and feedback, but weaken human capital when it becomes a substitute for problem solving. In politics and media, conversational systems can improve access to information while also making persuasion cheaper, more personalized, and more scalable.

At the same time, the picture for aggregate labor-market outcomes remains limited, and this is not necessarily inconsistent with the large micro-level gains reviewed in this paper. Productivity improvements on particular tasks do not mechanically translate into lower labor demand: firms may expand output, reallocate time toward complementary activities, or delay organizational adjustment. Survey evidence from U.S. executives is consistent with this interpretation, with firms reporting productivity gains but limited reductions in headcount so far.

It is also important to be clear about the boundaries of the current evidence. Most of the studies reviewed here concern chatbot-style systems: a human asks a question, the model responds, and the human retains control over what to do next. Much of the evidence is short-run, task-specific, and concentrated in relatively high-income or highly educated settings. The literature is much thinner on persistent general-equilibrium effects, public-sector outcomes, and low-income contexts, and it remains especially limited for more autonomous AI systems.

That last gap matters because the technology is already moving beyond chatbots toward agents that can decompose tasks, use tools, and act across multiple stages with more limited human oversight. The early work on such systems is still thin and mixed. The transition matters because it changes the economic question. When a system answers a query, the central issue is assis-

tance and productivity. When it begins to choose actions, the central issues become delegation, verification, and institutional design. Some mechanisms identified in the chatbot literature are likely to persist: lower information costs, larger gains for less experienced users, and the moral consequences of placing an intermediary between an actor and an outcome. Results that depend on a specific interface, or on a human reviewing each model output, are less likely to.

Across the domains reviewed here, the trade-offs that emerge follow a familiar shape: gains in competence and access have to be weighed against costs to oversight, learning, and accountability. In structured workplace tasks, productivity gains are large and oversight is relatively easy. In education, performance gains are real but the learning cost can dominate. In politics and public administration, the key concerns are whose interests the system serves and who is responsible when things go wrong. As systems become more autonomous, these trade-offs will become more central rather than less.

The next phase of empirical research should therefore move in two directions. One is outward: from short-run task performance to longer-run effects on wages, organizations, institutions, and social welfare. The other is forward: from chatbots that assist humans to agents that act with increasing autonomy. The first wave of evidence has shown that generative AI can be economically important. The harder question now is under what institutional and organizational conditions its benefits can be realized without eroding learning, accountability, and trust.

References

- [1] Acemoglu, Daron. 2025. "The Simple Macroeconomics of AI." *Economic Policy* 40(121): 13-58.
- [2] Acemoglu, Daron, and Pascual Restrepo. 2018. "The Race between Man and Machine: Implications of Technology for Growth, Factor Shares, and Employment." *American Economic Review* 108(6): 1488–1542.
- [3] Acemoglu, Daron, and Pascual Restrepo. 2019. "Automation and New Tasks: How Technology Displaces and Reinstates Labor." *Journal of Economic Perspectives* 33(2): 3–30.
- [4] Acemoglu, Daron, Asuman Ozdaglar, and James Siderius. 2026. "AI and Social Media: A Political Economy Perspective." In *The Political Economy of Artificial Intelligence*, edited by Ajay Agrawal, Joshua Gans, Avi Goldfarb, and Catherine E. Tucker. Chicago: University of Chicago Press.
- [5] Agrawal, Ajay K., John McHale, and Alexander Oettl. 2026. "AI in Science." NBER Working Paper 34953.
- [6] Alishani, Art, and Vincent Homburg. 2025. "When Citizens Meet the Chatbot: Evidence from a Survey Vignette Experiment in Estonia." *Public Policy and Administration*, published online, doi:10.1177/09520767251404286.
- [7] Althoff, Lukas, and Hugo Reichardt. 2026. "Task-Specific Technical Change and Comparative Advantage." CESifo Working Paper 12403.
- [8] Aoki, Naomi. 2020. "An Experimental Study of Public Trust in AI Chatbots in the Public Sector." *Government Information Quarterly* 37(4): 101490.
- [9] Aoki, Naomi. 2021. "The Importance of the Assurance that 'Humans Are Still in the Decision Loop' for Public Trust in Artificial Intelligence." *Computers in Human Behavior* 114: 106572.
- [10] Anthropic. 2026a. "Anthropic Economic Index Report: Economic Primitives." Research report, January.

- [11] Anthropic. 2026b. “What 81,000 People Want from AI.” Research report based on AI-led qualitative interviews with 80,508 users across 159 countries.
- [12] Ash, Elliott, Francesco Capozza, and Sergio Galletta. 2026. “Demand for Conversational AI.” CESifo Working Paper 12673.
- [13] Ash, Elliott, Sergio Galletta, and Giacomo Opocher. 2025. “BallotBot: Can AI Strengthen Democracy?” CEPR Discussion Paper 20070.
- [14] Bai, Hao, Jan G. Voelkel, Sarah Muldowney, Johannes C. Eichstaedt, and Robb Willer. 2025. “LLM-Generated Messages Can Persuade Humans on Policy Issues.” *Nature Communications* 16: 6037.
- [15] Baslandze, Salomé, Zachary Edwards, John Graham, Ty McClure, Brent H. Meyer, Michael Sparks, Sonya R. Waddell, and Daniel Weitz. 2026. “Artificial Intelligence, Productivity, and the Workforce: Evidence from Corporate Executives.” NBER Working Paper 34984.
- [16] Bastani, Hamsa, Osbert Bastani, Alp Sungu, Haosen Ge, Ozge Kabakcı, and Rei Mariman. 2025. “Generative AI without Guardrails Can Harm Learning: Evidence from High School Mathematics.” *Proceedings of the National Academy of Sciences* 122(26): e2422633122.
- [17] Bick, Alexander, Adam Blandin, and David J. Deming. 2026. “The Rapid Adoption of Generative AI.” *Management Science*, Articles in Advance. <https://doi.org/10.1287/mnsc.2025.02523>.
- [18] Bick, Alexander, Adam Blandin, and David J. Deming. 2025. “The State of Generative AI Adoption in 2025.” Federal Reserve Bank of St. Louis, *On the Economy*, November 13.
- [19] Biener, Christian, and Aline Waeber. 2024. “Would I Lie to You? How Interaction with Chatbots Induces Dishonesty.” *Journal of Behavioral and Experimental Economics* 112:102279.
- [20] Brynjolfsson, Erik, Bharat Chandar, and Ruyu Chen. 2025. “Canaries in the Coal Mine? Six Facts about the Recent Employment Effects of Artificial Intelligence.” Stanford Digital Economy Lab Working Paper.
- [21] Brynjolfsson, Erik, Danielle Li, and Lindsey Raymond. 2025. “Genera-

- tive AI at Work.” *Quarterly Journal of Economics* 140(2): 889–942.
- [22] Bursztyn, Leonardo, Alex Imas, Rafael Jiménez-Durán, Aaron Leonard, and Christopher Roth. 2026. "Social Dynamics of AI Adoption in Parents' Educational Decisions." *Proceedings of the National Academy of Sciences* 123(31): e2533193123.
- [23] Campante, Filipe R., Ruben Durante, Felix Hagemeister, and Ananya Sen. 2025. "GenAI Misinformation, Trust, and News Consumption: Evidence from a Field Experiment." NBER Working Paper 34100.
- [24] Canada Revenue Agency. 2025. "GenAI Chatbot Beta." Government of Canada webpage, updated November 7.
- [25] Costello, Thomas H., Gordon Pennycook, and David G. Rand. 2024. "Durably Reducing Conspiracy Beliefs through Dialogues with AI." *Science* 385(6714): eadq1814.
- [26] Cruces, Guillermo, Pablo Fernández Mejjide, Sebastián Galiani, Rubén Galvez, and Mariana Lombardi. 2026. "Does Generative AI Narrow Education-Based Productivity Gaps?" NBER Working Paper 34851.
- [27] Cui, Zheyuan, Mert Demirer, Sonia Jaffe, Leon Musolff, Sida Peng, and Tobias Salz. 2026. "The Effects of Generative AI on High-Skilled Work: Evidence from Three Field Experiments with Software Developers." *Management Science*, <https://doi.org/10.1287/mnsc.2025.00535>.
- [28] Chatterji, Aaron, Thomas Cunningham, David J. Deming, Zoe Hitzig, Christopher Ong, Carl Yan Shan, and Kevin Wadman. 2025. "How People Use ChatGPT." NBER Working Paper 34255.
- [29] Chen, Wilbur Xinyuan, Suraj Srinivasan, and Saleh Zakerinia. 2024. "Displacement or Complementarity? The Labor Market Impact of Generative AI." Harvard Business School Working Paper 25-039.
- [30] Chopra, Felix, Ingar Haaland, Nicolas Röver, and Christopher Roth. 2026. "Evaluating Behavioral Interventions at Scale with AI." CESifo Working Paper 12410.
- [31] Chung, Angel Tsai-Hsuan, Botong Zhang, Ling-Chieh Kung, Hamza Bastani, and Osbert Bastani. 2026. "Effective Personalized AI Tu-

- tors via LLM-Guided Reinforcement Learning.” SSRN Working Paper 6423358.
- [32] Dell’Acqua, Fabrizio, Edward McFowland III, Ethan Mollick, Hila Lifshitz-Assaf, Katherine C. Kellogg, Saran Rajendran, Lisa Kraye, François Cadelon, and Karim R. Lakhani. 2026a. “Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of AI on Knowledge Worker Productivity and Quality.” *Organization Science* 37(2): 403–423.
 - [33] Dell’Acqua, Fabrizio, Charles Ayoubi, Hila Lifshitz-Assaf, Raffaella Sadun, Ethan Mollick, Lilach Mollick, Yi Han, Jeff Goldman, Hari Nair, Stew Taub, and Karim R. Lakhani. 2026b. “The Cybernetic Teammate: A Field Experiment on Generative AI and Teamwork.” *Organization Science*, published online June 2026. <https://doi.org/10.1287/orsc.2025.20702>.
 - [34] De Freitas, Julian, Zeliha Oğuz-Uğuralp, Ahmet Kaan Uğuralp, and Stefano Puntoni. 2026. “AI Companions Reduce Loneliness.” *Journal of Consumer Research* 52(6): 1126–1148.
 - [35] DeVerna, Matthew R., Harry Yaojun Yan, Kai-Cheng Yang, and Filippo Menczer. 2024. “Fact-Checking Information from Large Language Models Can Decrease Headline Discernment.” *Proceedings of the National Academy of Sciences* 121(50): e2322823121.
 - [36] Dillon, Eleanor Wiske, Sonia Jaffe, Nicole Immorlica, and Christopher T. Stanton. 2025. “Shifting Work Patterns with Generative AI.” *American Economic Review: Insights*, forthcoming. <https://doi.org/10.1257/aeri.20250275>.
 - [37] Doshi, Anil R., and Oliver P. Hauser. 2024. “Generative AI Enhances Individual Creativity but Reduces the Collective Diversity of Novel Content.” *Science Advances* 10(28): adn5290.
 - [38] European Central Bank. 2026. “Artificial Intelligence: Friend or Foe for Hiring in Europe Today?” ECB Blog, March 4.
 - [39] Eurostat. 2025a. “20% of EU Enterprises Use AI Technologies.” Eu-

- rostat News, December 11.
- [40] Eurostat. 2025b. “32.7% of EU People Used Generative AI Tools in 2025.” Eurostat News, December 16.
 - [41] Eurostat. 2026. “64% of 16–24-Year-Olds Used AI in 2025.” Eurostat News, February 10.
 - [42] Fang, Cathy Mengying, Auren R. Liu, Valdemar Danry, Eunhae Lee, Samantha W.T. Chan, Pat Pataranutaporn, Pattie Maes, Jason Phang, Michael Lampe, Lama Ahmad, and Sandhini Agarwal. 2025. “How AI and Human Behaviors Shape Psychosocial Effects of Extended Chatbot Use: A Longitudinal Randomized Controlled Study.” arXiv:2503.17473.
 - [43] Facius, Dennis, and Roberto Iacono. 2026. “Labor Market Consequences of Generative AI: Early Evidence from Norway.” CESifo Working Paper 12752.
 - [44] Fischer, Mira, Holger A. Rau, and Rainer Michael Rilke. 2025. “AI Tutoring Enhances Student Learning Without Crowding Out Reading Effort.” IZA Discussion Paper 18338.
 - [45] Gallup. 2026. “Frequent Use of AI in the Workplace Continued to Rise in Q4.” Gallup Workplace, January 25.
 - [46] GOV.UK. 2025. “DSIT: GOV.UK Chat.” Algorithmic Transparency Record, October 7.
 - [47] gov.gr. 2026. “The Digital Assistant for gov.gr.” gov.gr webpage.
 - [48] Guarascio, Dario, and Jelena Reljic. 2025. “AI and Employment in Europe.” *Economics Letters* 247: 112183.
 - [49] Hackenburg, Kobi, Ben M. Tappin, Luke Hewitt, Ed Saunders, Sid Black, Hause Lin, Catherine Fist, Helen Margetts, David G. Rand, and Christopher Summerfield. 2025. “The Levers of Political Persuasion with Conversational Artificial Intelligence.” *Science* 390(6777): eaea3884.
 - [50] Hao, Qian Yue, Fengli Xu, Yong Li, and James A. Evans. 2026. “Artificial Intelligence Tools Expand Scientists’ Impact but Contract Science’s Focus.” *Nature* 649: 1237–1243.
 - [51] Hausman, Naomi, Oren Rigbi, and Sarit Weisburd. 2025. “Generative

- AI's Impact on Student Achievement and Implications for Worker Productivity." SSRN Working Paper 5393516.
- [52] Hui, Xiang, Oren Reshef, and Luofeng Zhou. 2024. "The Short-Term Effects of Generative Artificial Intelligence on Employment: Evidence from an Online Labor Market." *Organization Science* 35(6): 1977–1989.
- [53] Humlum, Anders, and Emilie Vestergaard. 2025. "Large Language Models, Small Labor Market Effects." NBER Working Paper 33777.
- [54] Jabarian, Brian, and Luca Henkel. 2025. "Voice AI in Firms: A Natural Field Experiment on Automated Job Interviews." SSRN Working Paper 5395709, August 18.
- [55] Johnston, Andrew C., and Christos A. Makridis. 2025. "AI, Output, and Employment." SSRN Working Paper 6460358.
- [56] Ju, Harang, and Sinan Aral. 2025. "Collaborating with AI Agents: Field Experiments on Teamwork, Productivity, and Performance." arXiv:2503.18238.
- [57] Kestin, Gregory, Kelly Miller, Anna Klaes, Timothy Milbourne, and Gregorio Ponti. 2025. "AI Tutoring Outperforms In-Class Active Learning: An RCT Introducing a Novel Research-Based Design in an Authentic Educational Setting." *Scientific Reports* 15: 17458.
- [58] Kauhanen, Antti, and Petri Rouvinen. 2026. "AI Has Not Impacted the Youth Labor Market in Finland." ETLA Working Paper 135.
- [59] Köbis, Nils, Zoe Rahwan, Raluca Rilla, Bramantyo Ibnu Supriyatno, Clara Bersch, Tamer Ajaj, Jean-François Bonnefon, and Iyad Rahwan. 2025. "Delegation to Artificial Intelligence Can Increase Dishonest Behaviour." *Nature* 646: 126–134.
- [60] Larsen, Anna Grøndahl, and Asbjørn Følstad. 2024. "The Impact of Chatbots on Public Service Provision: A Qualitative Interview Study with Citizens and Public Service Providers." *Government Information Quarterly* 41(2): 101927.
- [61] Liang, Zhehao, You Li, and Tao Chen. 2025. "When AI Fails, Who Gets the Blame? Citizens' Attribution Patterns in AI-Induced Public Service

- Failures” *Journal of Chinese Governance* 10(4): 630–659.
- [62] Lin, Hause, Gabriela Czarnek, Benjamin Lewis, Joshua P. White, Adam J. Berinsky, Thomas Costello, Gordon Pennycook, and David G. Rand. 2025. “Persuading Voters Using Human–Artificial Intelligence Dialogues.” *Nature* 648: 394–401.
- [63] Noy, Shakked, and Whitney Zhang. 2023. “Experimental Evidence on the Productivity Effects of Generative Artificial Intelligence.” *Science* 381(6654): 187–192.
- [64] OECD. 2025. “mAigov.” OECD AI Policy Observatory dashboard entry, July 9.
- [65] OECD. 2026. “The Agentic AI Landscape and Its Conceptual Foundations.” *OECD Artificial Intelligence Papers* 56.
- [66] OpenAI. 2025. “1 Million Business Customers.” OpenAI webpage, November 5.
- [67] Otis, Nicholas G., Rowan Clarke, Solène Delecourt, David Holtz, and Rembrand Koning. 2026. “The Uneven Impact of Generative Artificial Intelligence on Entrepreneurial Performance: Evidence from a Field Experiment in Kenya.” *Management Science*, Articles in Advance. <https://doi.org/10.1287/mnsc.2024.06909>.
- [68] Piccardi, Tiziano, Martin Saveski, Chenyan Jia, Jeffrey T. Hancock, Jenny L. Tsai, and Michael S. Bernstein. 2025. “Reranking Partisan Animosity in Algorithmic Social Media Feeds Alters Affective Polarization.” *Science* 390(6776): eadu5584.
- [69] gov.pt. 2024. “You Can Now Talk to the gov.pt Portal’s New Virtual Assistant.” gov.pt news release, December 27.
- [70] Salvi, Francesco, Manoel Horta Ribeiro, Riccardo Gallotti, and Robert West. 2025. “On the Conversational Persuasiveness of GPT-4.” *Nature Human Behaviour* 9(8): 1645–1653.
- [71] Strömberg, David, Victor Lei, and Yanhui Wu. 2026. “The Generative AI Learning Penalty: Evidence from Chinese Secondary Education.” CEPR Discussion Paper 21577.

- [72] Sarkar, Suproteem K. 2026. “AI Agents and Higher-Order Work.” SSRN Working Paper 5713646.
- [73] Wang, Rose E., Ana T. Ribeiro, Carly D. Robinson, Susanna Loeb, and Dorottya Demszky. 2024. “Tutor CoPilot: A Human-AI Approach for Scaling Real-Time Expertise.” EdWorkingPaper 24-1054, Annenberg Institute, Brown University.
- [74] Vanzo, Alessandro, Sankalan Pal Chowdhury, and Mrinmaya Sachan. 2025. “GPT-4 as a Homework Tutor Can Improve Student Engagement and Learning Outcomes.” In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 31119–31136.
- [75] Vaswani, Ashish, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. “Attention Is All You Need.” In *Advances in Neural Information Processing Systems* 30: 5998–6008.
- [76] Veale, Michael, Isra Black, Mark Dsouza, Matt Fisher, Mujtaba Ghaus, Thea Gibbs, Orla Lynskey, Colm O’Cinneide, Eloise Scotford, Olga Thomas, Alina Trapova, and Kimberley Trapp. 2025. “Artificial Intelligence, Education and Assessment at UCL Laws: Current Thinking and Next Steps for the UK Legal Education Sector.” UCL Legal Studies Research Paper 04/2025, SSRN Working Paper 5241868.
- [77] Velez, Yamil R., Donald P. Green, and Semra Sevi. 2025. “Chatbot Voting Advice Applications Inform but Seldom Sway Young Unaligned Voters.” *Proceedings of the National Academy of Sciences* 122(50): e2515516122.

Appendix

Study	Setting	Intervention	Design	Outcome	Horizon	Main result
<i>Work, productivity, and entrepreneurship</i>						
Brynjolfsson, Li, and Raymond (2025)	Customer support, large U.S. firm	Generative-AI assistant	Staggered rollout, DiD	Issues/hour	Up to 1 yr	+15% on average; largest gains for less experienced workers
Noy and Zhang (2023)	Online writing tasks	Access to ChatGPT	Online RCT	Time, quality	Single task	Time -40%, quality +18%
Dell’Acqua et al. (2026a)	BCG consultants	GPT-4 access on professional tasks	Field experiment	Quality, speed	Single session	Gains inside frontier; losses outside it
Cui et al. (2026)	Software developers, 3 firms	GitHub Copilot access	RCT, IV for usage	Tasks completed	Months	+26% weekly tasks; larger gains for less experienced devs
Dillon et al. (2025)	Knowledge workers, 66 firms	GenAI in email/docs/meetings	6-month RCT	Time use	6 months	Email time falls; faster doc completion; no large task-composition shift
Jabarian and Henkel (2025)	Approximately 70k job applicants	AI vs. human voice interviewers	Natural field experiment	Offers, retention	30 days	Higher offers, starts, and retention with AI interviews
Otis et al. (2026)	Kenyan small businesses	GPT-4 business assistant via WhatsApp	5-month RCT	Profit, sales	5 months	Zero ATE; -8% for low performers, +15% for high performers
<i>Macro and aggregate labor-market effects</i>						
Humlum and Vestergaard (2025)	Denmark, all workers	ChatGPT adoption + admin records	DiD on adoption + admin data	Earnings, hours	2 yrs	Near-zero effects; CI rules out >2%
Facius and Iacono (2026)	Norway, all workers, age cohorts	ChatGPT exposure shock	DiD, SDID, shift-share IV	Employment	Approximately 2.5 yrs	No robust displacement; estimates absorbed by pre-trends
Chen, Srinivasan, and Zakerinia (2024)	U.S. job postings	Occupation-level AI exposure	DiD by exposure type	Vacancies	Years	Demand falls in automating uses, rises in augmenting uses

Study	Setting	Intervention	Design	Outcome	Horizon	Main result
<i>Learning and human capital</i>						
Fischer, Rau, and Rilke (2025)	University students, exam prep	AI tutor access, restricted vs. open	RCT	Unaided test score	Hours–days	+0.23 SD overall; +0.34 SD with open access
Bastani et al. (2025)	Turkish high-school math	Raw GPT vs. guardrailed tutor	Field experiment	Unaided exam	Weeks	Raw GPT hurts later unaided performance; guardrails neutral
Strömberg, Lei, and Wu (2026)	Chinese secondary schools, grades 7–12	AI adoption, observed	30-month panel, 9 subjects	Homework, exams	6 mo.–2 yrs	Homework +18%; monthly exams -20%; entrance -18 to -24%
Wang et al. (2024)	U.S. K–12 underserved tutoring	Tutor-facing LLM copilot	RCT, 900 tutors	Topic mastery	Term	+9% mastery, +25% for students of lower-rated tutors
Kestin et al. (2025)	Harvard physics	AI tutor vs. active-learning lecture	Within-subjects crossover	Learning gain	Single session	AI tutor doubles learning per unit time
Chung et al. (2026)	Taipei high schools, programming	Adaptive RL-based AI tutor	5-month RCT, 770 students	Unassisted final exam	5 months	+0.15 SD, approximately 6–9 months of schooling
Hausman, Rigbi, and Weisburd (2025)	Israeli university	AI availability across courses	DiD across AI-compatible courses	Grades, downstream performance	2 yrs	+1 to +1.5 pts in AI-compatible courses; lower later non-AI performance

Study	Setting	Intervention	Design	Outcome	Horizon	Main result
<i>Politics, media, and the public sector</i>						
Ash, Galletta, and Opocher (2025)	California voters, 2024 ballot propositions	BailotBot, RAG over voter guide	Pre-registered 3-wave RCT	Knowledge, information costs	Weeks	+18% accuracy; lower perceived information costs; no learning carry-over
Bai et al. (2025)	U.S. adults	LLM- vs. human-written persuasion	Pre-registered survey experiment	Policy attitudes	Survey	2–4 pt shift; LLM persuades via facts, humans via authenticity
Salvi et al. (2025)	Online debaters	GPT-4 with personalization	Pre-registered 2 x 2 x 3 RCT	Opinion change	Single debate	GPT-4 with profile wins 64%; +81% odds of shifting opponent
Lin et al. (2025)	U.S., Canada, Poland elections	AI dialogue advocating a candidate	Pre-registered field experiments	Candidate preference	Days–weeks	Significant shifts; right-leaning models more factually inaccurate
Costello, Pennycook, and Rand (2024)	U.S. conspiracy believers	GPT-4 dialogue	Pre-registered RCT	Belief in conspiracy	2 months	-20% belief, persists
Piccardi et al. (2025)	X/Twitter users	LLM downrank/uprank of animosity	10-day field experiment	Affective polarization	10 days	Downranking lowers polarization by 2 pts on 100-pt scale
DeVerna et al. (2024)	U.S. adults	LLM fact-checks of headlines	Pre-registered experiment	Headline discernment	Survey	Discernment falls; users over-trust LLM verdicts
Campante, Durante, Hagemeister, and Sen (2025)	German newspaper readers	Salience of GenAI misinformation	Field experiment	Trust, subscription	Months	Lower trust; higher visits and retention to credible outlet

Study	Setting	Intervention	Design	Outcome	Horizon	Main result
<i>Innovation, science, and the production of ideas</i>						
Doshi and Hauser (2024)	Online writers, 293 participants	AI-assisted short stories	RCT	Creativity, diversity	Single task	Individual creativity up, collective diversity down
Hao et al. (2026)	Large-scale scientific-publication data	AI-tool adoption	Bibliometric DID	Impact, breadth	Years	Impact rises; topical focus narrows
<i>Chatbots and social behavior</i>						
Chopra et al. (2026)	Social media users	AI motivational-interviewing chatbot	RCT	Reported usage change	2+ weeks	Sustained reduction in self-reported social media use
Fang et al. (2025)	Adult chatbot users	4-week chatbot interaction	RCT	Loneliness, socialization	4 weeks	Higher daily use is associated with loneliness and dependence
Köbis et al. (2025)	Lab + online	Human vs. AI delegation of dishonest task	13 experiments / 4 studies	Dishonesty	Single task	AI delegation increases dishonest behaviour

Notes: The table summarizes a selected subset of the studies reviewed in the paper. The aim is not to be exhaustive but to provide a compact view of the empirical landscape across domains. For each entry the table reports the study, the setting, the intervention, the research design, the main outcome, the time horizon, and a one-line summary of the main result. Studies are grouped by domain.

PARTNER ISTITUZIONALI

CeSPEM

CENTRO STUDI DI
POLITICA ECONOMICA E MONETARIA
"MARIO ARCELLI"



UNIVERSITÀ
CATTOLICA
del Sacro Cuore

 **LUISS**

CASMEF Centro Arcelli
per gli Studi Monetari e Finanziari

æb
ASSONEBB



FONDAZIONE
DI PIACENZA E VIGEVANO



CONFCOMMERCIO
IMPRESE PER L'ITALIA

BUSINESS PARTNER

IMI
CORPORATE &
INVESTMENT
BANKING
INTESA  **SANPAOLO**

sara 
ti assicura

Deloitte.

SOSTENITORI

3D Works

Allianz Bank F.A.

Banca d'Italia

Banca Popolare del Cassinate

Banco Posta SGR

Confindustria Piacenza

CONSOB

EFPA - Italia

Fondazione Ave Verum

Fondazione Mario Arcelli

ISTAT

Istituto per il Credito
Sportivo e Culturale

TIM

Unicredit

Ver Capital SGR

Per rinnovare o attivare un nuovo abbonamento
effettuare un **versamento** su:

c/c bancario n. 10187 Intesa Sanpaolo
Via Vittorio Veneto 108/b- 00187 ROMA
IBAN IT92 M030 6905 0361 0000 0010 187

intestato a: **Editrice Minerva Bancaria s.r.l.**

oppure inviare una **richiesta** a:

amministrazione@editriceminervabancaria.it

Condizioni di abbonamento ordinario per il 2027

	Rivista Bancaria Minerva Bancaria bimestrale	Economia Italiana quadrimestrale	Rivista Bancaria Minerva Bancaria + Economia Italiana
Canone Annuo Italia (<i>print</i>)	€ 130,00	€ 100,00	€ 180,00
Canone Annuo Estero (<i>print</i>)	€ 185,00	€ 130,00	€ 260,00
Abbonamento WEB	€ 80,00	€ 70,00	€ 110,00
Canone Annuo Italia (<i>print + web</i>)	€ 170,00	€ 130,00	€ 260,00
Canone Annuo Estero (<i>print + web</i>)	€ 220,00	€ 160,00	€ 330,00

L'abbonamento è per un anno solare e dà diritto a tutti i numeri usciti nell'anno.

L'Amministrazione non risponde degli eventuali disguidi postali.

I fascicoli non pervenuti dovranno essere richiesti alla pubblicazione del fascicolo successivo.

Decorso tale termine, i fascicoli disponibili saranno inviati contro rimessa del prezzo di copertina.

Prezzo del fascicolo in corso **€ 50,00 / € 10,00** digitale

Prezzo di un fascicolo arretrato (annata precedente) **€ 60,00 / € 10,00** digitale

Pubblicità

1 pagina **€ 1.500,00** - 1/2 pagina **€ 800,00**

Editrice Minerva Bancaria
COMITATO EDITORIALE STRATEGICO

PRESIDENTE

GIORGIO DI GIORGIO, Luiss Guido Carli

COMITATO

CLAUDIO CHIACCHIERINI, Università degli Studi di Milano Bicocca

MARIO COMANA, Luiss Guido Carli

ADRIANO DE MAIO, Università Link Campus

RAFFAELE LENER, Università degli Studi di Roma Tor Vergata

MARCELLO MARTINEZ, Università della Campania

GIOVANNI PARRILLO, Editrice Minerva Bancaria

MARCO TOFANELLI, Assoreti

ECONOMIA ITALIANA 2026/2

Intelligenza artificiale: da innovazione tecnologica a crescita economica inclusiva

Il volume, editor **Alessandra Bonfiglioli** (Università di Bergamo), **Gino Gancia** (Università di Milano Bicocca) e **Paola Profeta** (Università Bocconi), offre un'attualissima analisi degli effetti dell'Intelligenza Artificiale sul sistema economico.

Una nuova *general purpose technology* suscita previsioni estreme: un'accelerazione senza precedenti della crescita per alcuni, la sostituzione del lavoro umano per altri. Se l'IA sta già cambiando organizzazioni e compiti cognitivi, resta da capire, osservano gli editor, *"attraverso quali meccanismi questi cambiamenti possano tradursi in trasformazioni dell'intero sistema economico e sociale"*.

Il filo conduttore dei cinque saggi è una domanda: **quali condizioni consentono all'IA di trasformarsi da innovazione tecnologica in crescita economica inclusiva?** Tre le questioni in gioco: gli effetti dell'IA sulla produttività, la redistribuzione di opportunità e rischi tra lavoratori, famiglie e imprese, il ruolo delle politiche pubbliche.

Flavio Calvino e Antonio Ughi leggono la diffusione dell'IA come un processo sistemico, i cui benefici dipendono da un ecosistema favorevole all'innovazione. **Sergio Galletta** ricostruisce le prime evidenze empiriche sull'IA generativa: la questione non è più il potenziale trasformativo dell'IA, ma le condizioni per una diffusione efficace e socialmente desiderabile. L'indagine di **Massimo Anelli, Italo Colantone, Aina Gallego e Piero Stanig** sui lavoratori italiani più esposti all'IA generativa rileva insieme ottimismo, cautela e una crescente domanda di politiche pubbliche. **Emilio Colombo e Francesco Trentini** propongono la famiglia, anziché il singolo lavoratore, come unità di analisi degli effetti distributivi: un meccanismo di assicurazione informale. Infine **Mauro Caselli, Enrico Cristofaletti e Andrea Fracasso** mostrano che le imprese italiane che adottano l'IA investono di più anche in benessere organizzativo, salute e sicurezza dei lavoratori. Ne emerge un'immagine dell'IA diversa da quella del dibattito pubblico: *"non una tecnologia che produce automaticamente vincitori e vinti"*, ma un insieme di innovazioni il cui impatto dipende dall'interazione tra progresso tecnico, organizzazione economica e istituzioni. Di qui le tre indicazioni di policy degli editor: **investimenti complementari, gestione della transizione occupazionale, governance dell'innovazione**. Come nelle rivoluzioni tecnologiche del passato, tutto dipenderà dalla capacità di imprese, lavoratori e istituzioni di riorganizzarsi.

Completano il numero l'intervento di **Carmine Di Noia**, con un'ampia ricerca (450 partecipanti, Authority incluse) sull'uso dell'IA nella finanza italiana, e la rubrica di **Massimiliano Parco** su sicurezza energetica, materie prime e nuove dipendenze.

ECONOMIA ITALIANA nasce nel 1979 per approfondire e allargare il dibattito sui nodi strutturali e i problemi dell'economia italiana, anche al fine di elaborare adeguate proposte strategiche e di *policy*. L'Editrice Minerva Bancaria è impegnata a portare avanti questa sfida e a fare di Economia Italiana il più vivace e aperto strumento di dialogo e riflessione tra accademici, *policy makers* ed esponenti di rilievo dei diversi settori produttivi del Paese.